

Stepping Lazy Programs

Stephen Chang
Northeastern University

Eli Barzilay
Northeastern University

John Clements
California Polytechnic State University

Matthias Felleisen
Northeastern University

Communicating Author: stchang@ccs.neu.edu

Abstract

Debugging lazy functional programs poses serious challenges. In support of the “stop, examine, and resume” debugging style of imperative languages, some debugging tools abandon lazy evaluation. Other debuggers preserve laziness but present it in a way that may confuse programmers because the focus of evaluation jumps around in a seemingly random manner.

In this paper, we introduce a supplemental tool, the algebraic program stepper. An algebraic stepper shows computation as a mathematical calculation. Algebraic stepping could be particularly useful for novice programmers or programmers new to lazy programming. Mathematically speaking, an algebraic stepper renders computation as the standard rewriting sequence of a lazy λ -calculus. Our novel lazy semantics introduces lazy evaluation as a form of parallel program rewriting. It represents a compromise between Launchbury’s store-based semantics and a simple, axiomatic description of lazy computation as sharing-via-parameters, à la Ariola et al. Finally, we prove that the stepper’s run-time machinery correctly reconstructs the standard rewriting sequence.

Categories and Subject Descriptors D.3.1 [Formal Definitions and Theory]: Semantics; D.3.2 [Language Classification]: Applicative (functional) languages; D.3.3 [Processors]: Debuggers

General Terms lazy programming; debugging and stepping; lazy lambda calculus

1. How Functional Programming Works

Hughes (1989) explains why functional programming matters. By “functional programming” Hughes specifically means “lazy” functional programming, and by “matters” he refers to the distinct advantages of laziness for programming with reusable components, i.e., functions and programs. Hughes’s examples demonstrate the ease of creating functions and programs by gluing together existing functions and programs. The key advantage of laziness is that even if one component appears to produce too much data, the laziness of the consuming component almost always naturally eliminates the superfluous data production. Since then, numerous researchers have extended this argument with examples of their own, usually accompanied by elegant code in Haskell.

Unfortunately, laziness also increases the distance between the programmer and the underlying machinery. Specifically, laziness reduces a programmer’s ability to predict when certain expressions

are evaluated during program execution. As long as things work, this cognitive dissonance poses no problems. When a program exhibits erroneous behavior, however, programmers are often at a loss. A programmer can turn to a debugger for help, but the evaluation of lazy programs is often confusing enough that lazy debuggers resort to hiding laziness from the programmer in order to display useful information (Wallace et al. 2001; Ennals and Peyton Jones 2003a; Allwood et al. 2009).

An ideal debugger should not modify the execution model of a program. The maintainers of the Glasgow Haskell Compiler (Peyton Jones et al. 1992) share this sentiment, since the bundled GHCi debugger abides by this ideal. The authors of the GHCi debugger (Marlow et al. 2007) state that their debugger “lets the programmer see the effects of laziness,” and therefore, “shows the programmer what is actually happening in their program at run-time”. Unfortunately, the authors also acknowledge that their debugger presents lazy computations in a way that is difficult to follow, mostly through seemingly random jumps from one place to another in the program.

To mitigate the drawbacks of existing lazy debuggers, we propose a supplementary tool, the algebraic stepper. PLT’s DrRacket (Findler et al. 2002) comes with such a stepper for the call-by-value Racket¹ language. Given a functional program, the algebraic stepper displays its standard reduction sequence in the by-value λ -calculus (Clements et al. 2001). Our experience suggests that this kind of tool especially benefits novice programmers when they try to understand small programs, and programmers who are new to the language and are trying to explore some linguistic feature.

DrRacket’s stepper presents evaluation of a program directly as a manipulation of the source, similar to the calculated manipulations of a student of mathematics. We conjecture that this model is suitable for addressing the confusing nature of lazy evaluation. In this paper, we present: (1) our stepper for Lazy Racket (Barzilay and Clements 2005); (2) its underlying semantics, a novel lazy λ -calculus; (3) and a proof that the stepper correctly implements the standard rewriting semantics of the calculus. Roughly speaking, Lazy Racket is a call-by-need language that uses the same evaluation mechanism as Haskell.² A Lazy Racket module macro-expands its surface syntax into a plain Racket program enriched with appropriate `delay` and `force` constructs (Hatcliff and Danvy 1997). Lazy Racket is mostly used in educational settings—where we have tested the prototype of the stepper so far—though some programmers have found it useful to construct parser combinators or game trees in Lazy Racket modules, and then to export such pieces to plain Racket modules.

The novel lazy λ -calculus introduces the idea of selective parallel reduction to simulate shared reductions. On the one hand, it is

¹DrRacket and Racket were formerly known as DrScheme and PLT Scheme, respectively.

²The implementation of Lazy Racket is similar to SRFI 45 of the Scheme Standard.

nearly trivial to prove an equivalence between the standard rewriting semantics of our calculus and a Launchbury-style, store-based semantics (Launchbury 1993). On the other hand, the calculus is the appropriate basis for a correctness proof of the stepper. For the correctness proof, we construct a model of the `delay-and-force` implementation, further enriched with continuation marks (Clements et al. 2001), show that it bisimulates the standard rewriting semantics, and finally, exploit Clements's strategy for the rest of the proof (Clements 2006).

Section 2 briefly introduces Lazy Racket and its stepper. Section 3 presents a novel lazy λ -calculus and its essential theorems. Section 4 summarizes the essential details of the implementation of Lazy Racket and presents a model of the lazy stepper. Section 5 presents a correctness proof for the stepper and the penultimate section summarizes related work.

2. Lazy Racket and Its Stepper

Lazy Racket programs are sequences of `defines` and expressions that usually refer to the definitions. Here is a basic example:

```
(define (f x) (+ x x))
(f (+ 1 (+ 2 3)))
```

A programmer invokes the Lazy Racket stepper from the DrRacket IDE. Running the stepper displays the reduction sequence for the current program. Figure 1 shows a sequence of screenshots stepping through the above program, with each shot displaying one reduction step. A green box highlights the redex(es) on the left-

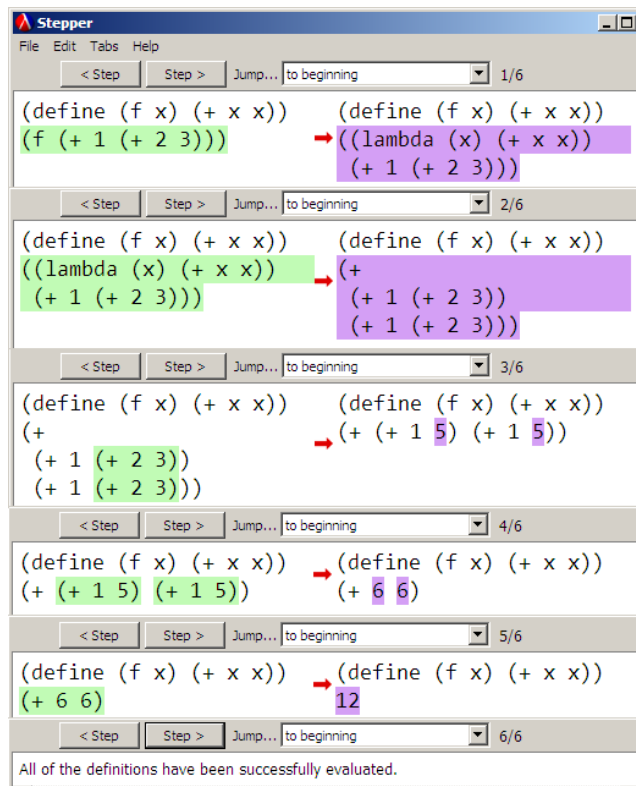


Figure 1. Lazy Stepper Example 1

hand side of a reduction step and a purple box highlights the contractum(s) on the right-hand side. The programmer can navigate the reduction sequence in either the forward or backward direction. Additional navigation features are in the planning stages.

In step 2, evaluation of the function argument is delayed so an unevaluated argument replaces each instance of the variable `x` in the function body. In step 3, evaluation of the program at the first `x` position requires the value of the argument, so the argument is forced in steps 3 and 4. In steps 3 and 4, all shared instances of the argument are reduced simultaneously. That is, *the stepper explains evaluation as an algebraic process using a form of parallel reduction*. Since the second `x` refers to the same delayed computation as the first `x`, by the time evaluation of the program requires a value at the second `x` position, a result is already available because the computed value of the first `x` was saved. In short, no argument evaluation is repeated.

A second example introduces the lazy `take!` function, which extracts the first `n` elements of a specified list:

```
(define (take! n lst)
  (if (= n 0)
      null
      (cons (first lst)
            (take! (- n 1) (rest lst)))))
(define (f lst) (+ (first lst) (second lst)))
(f (take! 3 (list 1 2 (/ 1 0) 4)))
```

The reduction sequence for this program, as viewed in the stepper, appears in figure 2. For space reasons, only interesting steps are shown.

In this example, the result of the `take!` computation is the argument to the function `f`. The `take!` computation extracts the first three elements of its list argument, but `f` only uses the first two list elements, so the third element, which produces an error, should not be forced. In step 3, the `take!` computation is forced because both `+` and `first` are strict. In Lazy Racket, `cons` behaves lazily and does not evaluate its arguments (Friedman and Wise 1976), so in step 5, the result of the `take!` computation is a `cons` with two thunks: one that retrieves the first element of the list, and one that contains the next iteration of the `take!` computation. In step 7, the first addition operand is finally reduced to a value. Notice that the first element in the argument to `second` is already reduced as well. The remaining steps force the next iteration of `take!` and similarly extract the second element of the list. Since only the first two elements of the list are needed, no additional `take!` iterations are computed and the division by zero never raises an error.

A third example involves infinite lists:

```
(define (add-one x) (+ x 1))
(define nats (cons 1 (map add-one nats)))
(+ (second nats) (third nats))
```

More importantly, it involves `map`, which the stepper has not annotated because it is a library function. The reduction sequence for this program appears in figure 3. Again, some function definitions and reduction steps have been elided from the screenshots. In step 1, the evaluation of `second` forces the `map` expression to a `cons` containing two thunks. Unlike the second example, the thunks are displayed as `<DelayedEvaluation#1>` and `<DelayedEvaluation#2>` because their contents are unknown, i.e., they were not part of the source program. In step 2, the `second` expression extracts `<DelayedEvaluation#1>` from the list, but the thunk is still unevaluated. In steps 3 and 4, evaluation of the program requires the value of `<DelayedEvaluation#1>`, so it is forced. Observe how the stepper updates the `nats` definition with the result as well. The remaining steps show the similar evaluation of the other addition operand and are thus omitted.

As a final example, we use our stepper to understand the behavior of a program presented by Marlow et al. (2007):

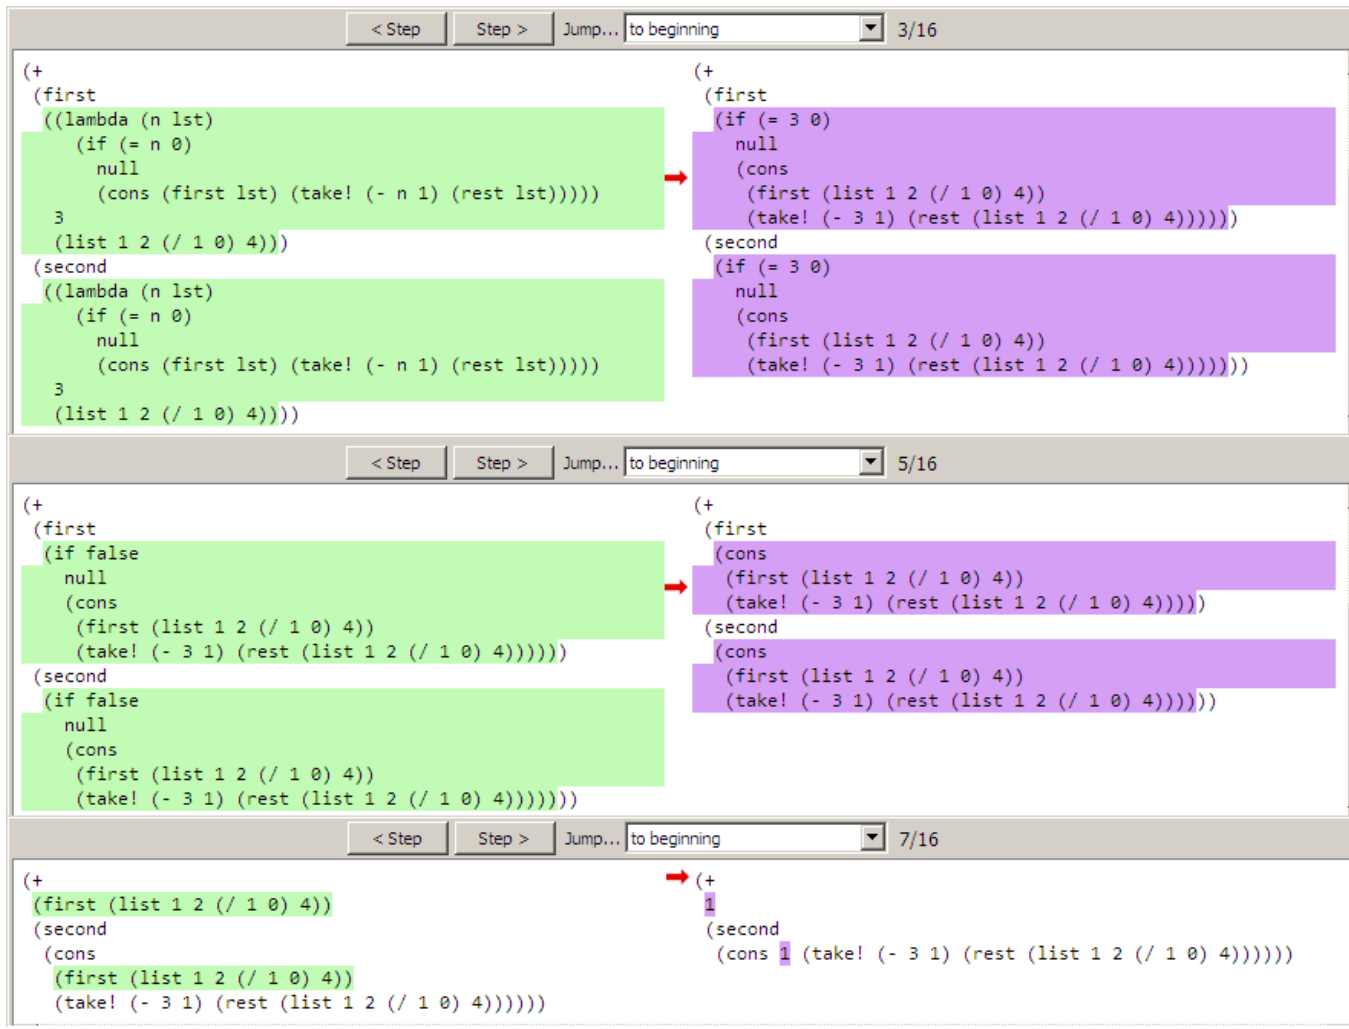


Figure 2. Lazy Stepper Example 2

```
;; [Listof Char] -> [Listof [Listof Char]]
(define (lines s)
  (cond
    [(null? s) null]
    [else
     (define-values (l t)
       (break (lambda (x) (equal? "\n" x)) s))
     (cons l (cond
              [(null? t) null]
              ; drop "\n" char and recur
              [else (lines (cdr t))]))]))

;; [Char -> Boolean] [Listof Char]
;; ->* [Listof Char] [Listof Char]
(define (break p? l)
  (let L ([l l] [line null])
    (cond
      [(null? l) (values (reverse line) null)]
      [(p? (car l)) (values (reverse line) l)]
      [else (L (cdr l) (cons (car l) line))]))
```

The `break` function consumes two arguments, a predicate on characters and a string represented as a list of characters, and splits the string at the first character for which the predicate is true, returning

two substrings simultaneously.³ The delimiting character remains as the first character of the second substring. The `lines` function uses `break` to separate an input string into lines, where a `"\n"` character begins a new line. Unlike `break`, the delimiting character is not included in the output of `lines`.

Evaluating the expression `(lines '("\n" "a"))` produces the expected result, an empty line (the empty string) and a line with one "a" character (the `!!` function is a recursive force function):

```
> (!! (lines '("\n" "a")))
'(() ("a"))
```

However, `(lines '("a" "\n"))` produces only one line:

```
> (!! (lines '("a" "\n")))
'(("a"))
```

Marlow et al. (2007) show how to use the GHCi debugger to understand this behavior. Figure 4 demonstrates that the stepper provides a superior vehicle in this situation. It specifically shows how `break` returns two values, which when evaluated, produce a one-character string "a" and a one-character string "\n", respectively. From the remainder of the definition of `lines`, we can deduce that the "a"

³The `values` Racket construct enables multiple return values.

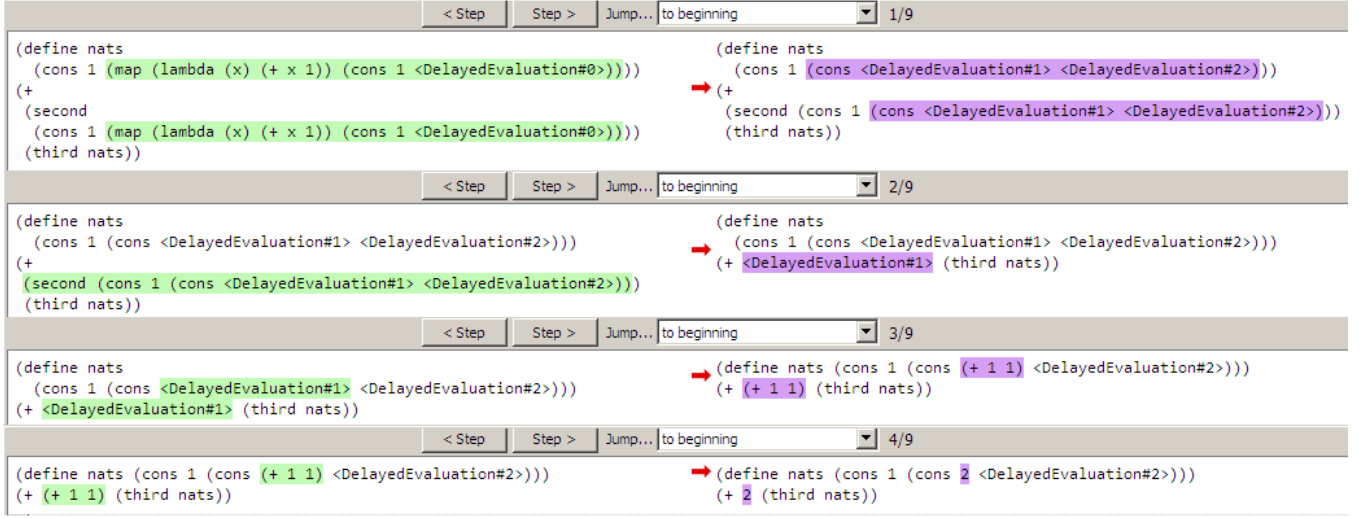


Figure 3. Lazy Stepper Example 3

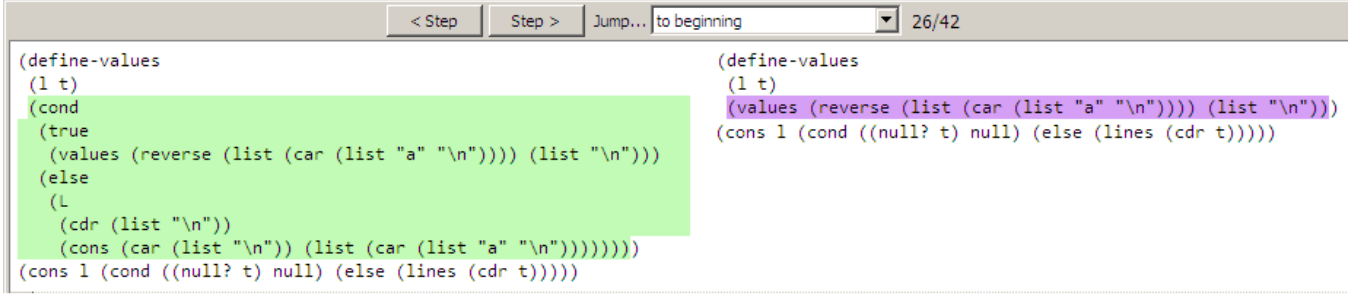


Figure 4. Lazy Stepper Example 4

string is retained, while the `"\n"` is dropped by the subsequent call to `cdr`, causing the recursive `lines` call to return an empty list, i.e. no lines, thus explaining the missing line. If we want `lines` to output a final, empty line when there is a trailing `"\n"` in the input, the base case must return a list with an empty line, `'()`, instead of just an empty list.

3. Lazy Racket Semantics

Our key theoretical innovation is the novel semantic view of laziness displayed in our stepper. Following tradition, we present this idea in the context of a λ -calculus, λ_{LR} :

$$\begin{aligned}
 e &= n \mid s \mid b \mid x \mid \lambda x. e \mid (e \ e) \mid (p^2 \ e \ e) \\
 &\mid (\text{cons } e \ e) \mid \text{null} \mid (p^1 \ e) \mid (\text{if } e \ e \ e) \\
 n &\in \mathbb{Z}, \ s \in \text{Strings}, \ b = \text{true} \mid \text{false} \\
 p^2 &= + \mid - \mid * \mid / \quad p^1 = \text{first} \mid \text{rest}
 \end{aligned}$$

The syntax of λ_{LR} is identical to the core of most functional programming languages and includes integers, strings, booleans, variables, abstractions, applications, primitives, lists, and a conditional.

To specify the semantics of λ_{LR} , we first extend e by adding a new expression:

$$e^{LR} = e \mid e^{LR\ell} \quad \ell \in \text{Labels}$$

The “labeled” expression, $e^{LR\ell}$, consists of a tag ℓ and a subexpression e^{LR} . Labeled expressions are not part of the language syn-

tax but are necessary for evaluation. Rewriting a labeled expression triggers the simultaneous rewriting of all other expressions that share the same label. Otherwise, labeled expressions do not affect program evaluation. The stepper renders labeled expressions without the label.

We require one constraint on labeled expressions in our language: all expressions with the same label ℓ must be identical. We call this the consistent labeling property:

Definition 1. A program is consistently labeled if, for all ℓ_1, ℓ_2 , e_1^{LR}, e_2^{LR} , if $e_1^{LR\ell_1}$ and $e_2^{LR\ell_2}$ are two subexpressions in a program, and $\ell_1 = \ell_2$, then $e_1^{LR} = e_2^{LR}$.

3.1 Rewriting Rules

To further formulate a semantics, we define the notion of values:

$$v = n \mid s \mid b \mid \lambda x. e^{LR} \mid \text{null} \mid (\text{cons } e^{LR\ell} \ e^{LR\ell}) \mid v^\ell$$

Numbers, strings, booleans, abstractions, and `null` are values. In addition, `cons` expressions where each element is labeled are also values. Finally, any value tagged with labels is also a value.

In the rewriting of λ_{LR} programs, evaluation contexts are used to determine which part of the program to rewrite next. Evaluation contexts are expressions where a hole `[]` replaces one subexpression:

$$\begin{aligned}
 E &= [] \mid (E \ e^{LR}) \mid (p^2 \ E \ e^{LR}) \mid (p^2 \ v \ E) \mid (p^1 \ E) \\
 &\mid (\text{if } E \ e^{LR} \ e^{LR}) \mid E^\ell
 \end{aligned}$$

LHS	\mapsto_{LR}	RHS	
	$\nexists E_1, E_2, \ell$ s.t. $E[] = E_1[(E_2[])^\ell]$ (redex not under label)	$E[] = E_1[(E_2[])^\ell]$ $\nexists E_3, E_4, \ell'$ s.t. $E_2[] = E_3[(E_4[])^{\ell'}]$ (redex occurs under label)	
$E[(\lambda x. e_1^{LR}) \vec{e}_2^{LR}]$	$E[e_1^{LR}\{x := e_2^{LR\ell_1}\}]$ fresh ℓ_1	$E[e_1^{LR}\{x := e_2^{LR\ell_1}\}]\{\ell \Leftarrow E_2[e_1^{LR}\{x := e_2^{LR\ell_1}\}]\}$ fresh ℓ_1	β_{LR}
$E[(p^2 n_1 \vec{e}_2)]$	$E[(\delta(p^2 n_1 n_2))]$	$E[(\delta(p^2 n_1 n_2))]\{\ell \Leftarrow E_2[(\delta(p^2 n_1 n_2))]\}$	PRIM
$E[(\text{cons } e_1^{LR} e_2^{LR})]$ e_1^{LR} unlabeled or e_2^{LR} unlabeled	$E[(\text{cons } e_1^{LR\ell_1} e_2^{LR\ell_2})]$ fresh ℓ_1, ℓ_2	$E[(\text{cons } e_1^{LR\ell_1} e_2^{LR\ell_2})]\{\ell \Leftarrow E_2[(\text{cons } e_1^{LR\ell_1} e_2^{LR\ell_2})]\}$ fresh ℓ_1, ℓ_2	CONS
$E[(\text{first } (\text{cons } e_1^{LR} e_2^{LR}))]$	$E[e_1^{LR}]$	$E[e_1^{LR}]\{\ell \Leftarrow E_2[e_1^{LR}]\}$	FIRST
$E[(\text{rest } (\text{cons } e_1^{LR} e_2^{LR}))]$	$E[e_2^{LR}]$	$E[e_2^{LR}]\{\ell \Leftarrow E_2[e_2^{LR}]\}$	REST
$E[(\text{if true } \vec{e}_1^{LR} e_2^{LR})]$	$E[e_1^{LR}]$	$E[e_1^{LR}]\{\ell \Leftarrow E_2[e_1^{LR}]\}$	IF-TRUE
$E[(\text{if false } \vec{e}_1^{LR} e_2^{LR})]$	$E[e_2^{LR}]$	$E[e_2^{LR}]\{\ell \Leftarrow E_2[e_2^{LR}]\}$	IF-FALSE

Figure 5. The λ_{LR} Reduction System.

The $(E e^{LR})$ context indicates that the operator in an application is evaluated before it is applied. The p^1 and p^2 contexts indicate that these primitives are strict in all argument positions. The if context dictates strict evaluation of only the test expression. Finally, the E^ℓ context dictates that a redex search goes under labeled expressions. Essentially, when searching for a redex, expressions tagged with a label are treated as if they were unlabeled.

Evaluation of a λ_{LR} program proceeds according to the program rewriting system in figure 5. For each possible rewriting step, the program in the first column is partitioned into a redex and a context, and is rewritten to either the program in the second or the third column. In the second and third columns, the redex is always contracted. If the redex does not occur under a label, then it is the only contracted part of the program (column two). If the redex occurs under a label, all other instances of the label are similarly contracted (column three). In the third column, the context is further subdivided as $E[] = E_1[(E_2[])^\ell]$ where ℓ is the label nearest the redex, E_1 is the context around the ℓ -labeled expression, and E_2 is the context under label ℓ surrounding the redex. Thus E_2 contains no additional labels. An “update” function is used to perform the parallel reduction. The update function uses the notation $e_1^{LR}\{\ell \Leftarrow e_2^{LR}\}$ to mean that all expressions in e_1^{LR} immediately under a label ℓ are replaced with e_2^{LR} . The function is formally defined in figure 6. The last clause in the definition covers all cases not included by the preceding clauses.

The β_{LR} rule specifies that function application occurs before the evaluation of arguments. To remember where expressions originate, the argument receives a label ℓ_1 before substitution is performed. The notation $e^{LR\vec{\ell}}$ represents an expression wrapped in one or more labels. During a rewriting step, labels are discarded from values because no further reduction is possible.

Binary primitive applications are strict in their arguments, as seen in the PRIM rule. The δ function interprets binary primitive applications and is defined in the standard way (division by 0 results in a stuck state).

The CONS rule shows that, if either argument is unlabeled, both arguments are wrapped with labels. Adding an extra label around an already labeled expression will not change the rewriting sequence of the program because the parallel updating function only uses the innermost label. The FIRST and REST rules extract the first and second components from a cons cell, respectively, and the IF-TRUE and IF-FALSE rules similarly choose the first or second branch of the if expression.

$$\begin{aligned}
e_1^{LR\ell}\{\ell \Leftarrow e_2^{LR}\} &= e_2^{LR\ell} \\
e_1^{LR\ell_1}\{\ell_2 \Leftarrow e_2^{LR}\} &= (e_1^{LR}\{\ell_2 \Leftarrow e_2^{LR}\})^{\ell_1}, \ell_1 \neq \ell_2 \\
(\lambda x. e_1^{LR})\{\ell \Leftarrow e_2^{LR}\} &= \lambda x. (e_1^{LR}\{\ell \Leftarrow e_2^{LR}\}) \\
(e_1^{LR} e_2^{LR})\{\ell \Leftarrow e_3^{LR}\} &= (e_1^{LR}\{\ell \Leftarrow e_3^{LR}\} \\
&\quad e_2^{LR}\{\ell \Leftarrow e_3^{LR}\}) \\
(p^2 e_1^{LR} e_2^{LR})\{\ell \Leftarrow e_3^{LR}\} &= (p^2 e_1^{LR}\{\ell \Leftarrow e_3^{LR}\} \\
&\quad e_2^{LR}\{\ell \Leftarrow e_3^{LR}\}) \\
(\text{cons } e_1^{LR} e_2^{LR})\{\ell \Leftarrow e_3^{LR}\} &= (\text{cons } e_1^{LR}\{\ell \Leftarrow e_3^{LR}\} \\
&\quad e_2^{LR}\{\ell \Leftarrow e_3^{LR}\}) \\
(p^1 e_1^{LR})\{\ell \Leftarrow e_2^{LR}\} &= (p^1 e_1^{LR}\{\ell \Leftarrow e_2^{LR}\}) \\
(\text{if } e_1^{LR} e_2^{LR} e_3^{LR})\{\ell \Leftarrow e_4^{LR}\} &= (\text{if } e_1^{LR}\{\ell \Leftarrow e_4^{LR}\} \\
&\quad e_2^{LR}\{\ell \Leftarrow e_4^{LR}\} \\
&\quad e_3^{LR}\{\ell \Leftarrow e_4^{LR}\}) \\
\text{otherwise, } e_1^{LR}\{\ell \Leftarrow e_2^{LR}\} &= e_1^{LR}
\end{aligned}$$

Figure 6. Definition of parallel update function.

Program rewriting preserves the consistent labeling property.

Lemma 1. *If e_1^{LR} is consistently labeled and $e_1^{LR} \mapsto_{LR} e_2^{LR}$, then e_2^{LR} is consistently labeled.*

The rewriting rules are deterministic because any expression e^{LR} can be uniquely partitioned into an evaluation context and a redex. Thus if e_1^{LR} rewrites to a expression e_2^{LR} , then e_1^{LR} rewrites to e_2^{LR} in a canonical manner.

We can then use \mapsto_{LR} to define an evaluator:

$$\text{eval}_{LR}(e) \begin{cases} v, & \text{if } e \mapsto_{LR} v \\ \perp, & \text{if, for all } e \mapsto_{LR} e_1^{LR}, e_1^{LR} \mapsto_{LR} e_2^{LR} \\ \text{error}, & \text{if } e \mapsto_{LR} e_1^{LR}, e_1^{LR} \notin v, \\ & \nexists e_2^{LR} \text{ such that } e_1^{LR} \mapsto_{LR} e_2^{LR} \end{cases}$$

where \mapsto_{LR}^* is the reflexive-transitive closure of \mapsto_{LR} .

3.2 A New Call-by-Need Lambda Calculus

As is, the \mapsto_{LR} relation cannot describe the standard reduction sequences of any calculus. Each rule in figure 5 replaces the *entire* program with another, in a non-compositional manner. Put differently, the table of relations does not show how a standard reduction semantics is created from a basic notion of reduction (Barendregt 1984), like β , made compatible with the syntactic constructions of a leftmost-outermost context (Felleisen et al. 2009; Plotkin 1975).

In this section, we sketch how our rewriting semantics relates to a plain, yet novel call-by-need calculus. Put concisely, the calculus replaces the *deref* notion of reduction in Ariola and Felleisen’s (1997) call-by-need calculus⁴ with a rule that exploits the function parameter for sharing but implements evaluation via substitution. The calculus comes with a standard reduction theorem, and it is possible to show that our rewriting semantics essentially relates the program to its answer via the same steps as the standard reduction sequences of the calculus. Due to a lack of space, we merely spell out the basic ideas without stating any theorems or proofs.

For simplicity, we take the syntax of the λ -calculus as the starting point:

$$e^{AF} = x \mid \lambda x. e^{AF} \mid (e^{AF} e^{AF})$$

Evaluation of a λ_{AF} program terminates when it is reduced to an answer a^{AF} :

$$\begin{aligned} v^{AF} &= \lambda x. e^{AF} \\ a^{AF} &= v^{AF} \mid ((\lambda x. a^{AF}) e^{AF}) \end{aligned}$$

Programs reduce to answers instead of values because reduction does not remove application terms. The specification of a notion of reduction relies on the notion of an evaluation context E^{AF} :

$$E^{AF} = [] \mid (E^{AF} e^{AF}) \mid ((\lambda x. E^{AF}[x]) E^{AF}) \mid ((\lambda x. E^{AF}) e^{AF})$$

The evaluation contexts, especially the third one, specify that arguments are not evaluated until they are needed by some variable in the function body.

Here are the three notions of reductions (axioms) from Ariola and Felleisen’s call-by-need calculus:

$$\begin{aligned} ((\lambda x. E^{AF}[x]) v^{AF}) &= ((\lambda x. E^{AF}[v^{AF}]) v^{AF}) & (deref) \\ (((\lambda x. a^{AF}) e^{AF}) e^{AF'}) &= ((\lambda x. (a^{AF} e^{AF'})) e^{AF}) & (lift) \end{aligned}$$

$$((\lambda x. E^{AF}[x]) ((\lambda y. a^{AF}) e^{AF})) = ((\lambda y. ((\lambda x. E^{AF}[x]) a^{AF})) e^{AF}) \quad (assoc)$$

The *deref* axiom substitutes the evaluated argument for the variable in the function body. The other two axioms deal with answers that may appear on the left-hand or right-hand side of an application.

One problem is that the *deref* axiom leaves the application alone, even after the argument has been reduced to a value. Clearly, doing so contradicts both the natural implementations (which use a mix of graph rewriting and memoizing) and the widely used Launchbury semantics (which uses a store-based semantics to mimic memoizing). To get closer to this semantics, we propose to replace the *deref* axiom with the following β_{need} axiom:

$$((\lambda x. E^{AF}[x]) v^{AF}) = E^{AF}[x]\{x := v^{AF}\} \quad (\beta_{need})$$

This axiom says that when a parameter occurs in a “demand” position—the hole of an evaluation context—and the argument is a value, then plain old substitution captures the essence of parameter passing in a call-by-need language.

The new axiom is indeed a proper notion of reduction and is applicable in any context. Like Ariola and Felleisen’s calculus, our

⁴The calculus of Maraist et al. (1998) is unrelated in this case.

revised lazy λ -calculus is confluent (satisfies the Church-Rosser property) and comes with a Curry-Feys style standard reduction theorem. Furthermore, there is a simple bisimulation that relates standard reduction sequences to the semantics of figure 5.

4. Lazy Stepper Implementation

Figure 7 summarizes the software architecture of our stepper. The first row depicts a λ_{LR} Lazy Racket rewriting sequence. To construct this rewriting sequence, the lazy stepper first macro-expands a Lazy Racket program into a functional Racket program, enriched with *delay* and *force*, as mentioned in section 1. In turn, the stepper for (eager) Racket annotates the expanded program. Executing an annotated Racket program emits a series of output values, from which the reduction sequence for the unannotated Racket program is reconstructed. Once the lazy stepper has the plain Racket reduction sequence, it synthesizes each step to assemble the desired Lazy Racket rewriting sequence.

The correctness of the lazy stepper thus depends on two claims:

1. The reduction sequence of a plain Racket program can be reconstructed from the output produced when evaluating an annotated version of that program.
2. The rewriting sequence of a Lazy Racket program is equivalent to the reduction sequence of the corresponding plain Racket program, modulo macro-expansion and synthesis steps.

The first point corresponds to the work of Clements et al. (2001) and is depicted by the bottom half of figure 7. The second point is depicted by the top half of figure 7. The rest of the section formally presents the architecture in enough detail so that our stepper can be implemented for other programming languages, and so that we can prove its correctness. The actual correctness theorem and proof can be found in the next section.

4.1 Racket + delay/force

When the stepper is invoked on a Lazy Racket program, the source is first macro-expanded to a plain Racket program. Racket programs are eagerly evaluated, so lazy evaluation is simulated with the insertion of *delay* and *force* constructs. We model this expanded language with λ_{DF} , a core calculus of functional Racket with *delay* and *force*:

$$\begin{aligned} e^{RKT} &= n \mid s \mid b \mid x \mid \lambda x. e^{RKT} \mid (e^{RKT} e^{RKT}) \mid (p^2 e^{RKT} e^{RKT}) \\ &\quad \mid (\text{cons } e^{RKT} e^{RKT}) \mid \text{null} \mid (p^1 e^{RKT}) \\ &\quad \mid (\text{if } e^{RKT} e^{RKT} e^{RKT}) \mid (\text{delay } e^{RKT}) \mid (\text{force } e^{RKT}) \\ n &\in \mathbb{Z}, \quad s \in \text{Strings}, \quad b = \text{true} \mid \text{false} \\ p^2 &= + \mid - \mid * \mid / \quad p^1 = \text{first} \mid \text{rest} \end{aligned}$$

The syntax of λ_{DF} is similar to λ_{LR} except that *delay* and *force* replace labeled expressions. The *delay* construct suspends evaluation of its argument in a thunk; applying *force* to a thunk evaluates the suspended computation and memoizes the result. In addition, applying *force* to a suspended computation wrapped in multiple, nested thunks forces all the thunks, while applying *force* to a value returns that value.

The semantics of λ_{DF} combines the usual call-by-value world with store effects. We describe it with a high-level abstract machine, specifically, a CS machine (Felleisen et al. 2009). The C in the CS machine stands for control string and the S is a store that represents physical memory. In our machine the control string is an expression that may contain locations, i.e., references to delayed expressions in the store. In contrast to the standard CS machine, the store in our machine only holds delayed computations.

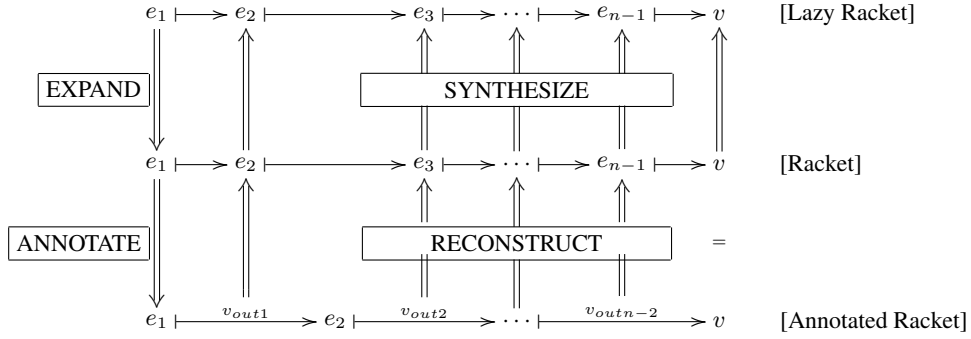


Figure 7. Stepper Implementation Architecture

\mapsto_{CS}			
$\langle E^{DF}[(\lambda x. e^{DF}) v^{DF}], \sigma \rangle$	\mapsto_{CS}	$\langle E^{DF}[e^{DF}\{x := v^{DF}\}], \sigma \rangle$	β_v
$\langle E^{DF}[(p^2 v_1^{DF} v_2^{DF})], \sigma \rangle$	\mapsto_{CS}	$\langle E^{DF}[(\delta (p^2 v_1^{DF} v_2^{DF}))], \sigma \rangle$	PRIM
$\langle E^{DF}[(\text{first } (\text{cons } v_1^{DF} v_2^{DF}))], \sigma \rangle$	\mapsto_{CS}	$\langle E^{DF}[v_1^{DF}], \sigma \rangle$	FIRST
$\langle E^{DF}[(\text{rest } (\text{cons } v_1^{DF} v_2^{DF}))], \sigma \rangle$	\mapsto_{CS}	$\langle E^{DF}[v_2^{DF}], \sigma \rangle$	REST
$\langle E^{DF}[(\text{if true } e_1^{DF} e_2^{DF})], \sigma \rangle$	\mapsto_{CS}	$\langle E^{DF}[e_1^{DF}], \sigma \rangle$	IF-TRUE
$\langle E^{DF}[(\text{if false } e_1^{DF} e_2^{DF})], \sigma \rangle$	\mapsto_{CS}	$\langle E^{DF}[e_2^{DF}], \sigma \rangle$	IF-FALSE
$\langle E^{DF}[(\text{delay } e^{DF})], \sigma \rangle$	\mapsto_{CS}	$\langle E^{DF}[\ell], \sigma[\ell \leftarrow e^{DF}] \rangle$	DELAY
		$\ell \notin \text{dom}(\sigma)$	
$\langle E^{DF}[(\text{force } \ell)], \sigma \rangle$	\mapsto_{CS}	$\langle E^{DF}[(\text{force } (\text{force } \ell \sigma[\ell]))], \sigma \rangle$	FORCE-DELAY
$\langle E^{DF}[(\text{force } \ell v^{DF})], \sigma \rangle$	\mapsto_{CS}	$\langle E[v^{DF}], \sigma[\ell \leftarrow v^{DF}] \rangle$	FORCE-UPDATE
$\langle E^{DF}[(\text{force } v^{DF})], \sigma \rangle$	\mapsto_{CS}	$\langle E^{DF}[v^{DF}], \sigma \rangle$	FORCE-NONDELAY
$v^{DF} \notin \text{Locations}$			

Figure 8. CS Machine Transitions

Here is the specification of our CS machine:

$\mathcal{P}^{DF} = \langle c^{DF}, \sigma \rangle$	(Machine States)
$\mathcal{S}^{DF} = \mathcal{P}_1^{DF}, \dots, \mathcal{P}_n^{DF}$	(Transition Sequences)
$c^{DF} = E^{DF}[e^{DF}]$	(Control Strings)
$e^{DF} = e^{RKT} \mid \ell$	(Machine Expressions)
$\sigma = ((\ell, e^{DF}), \dots)$	(Stores)
$\ell \in \text{Locations}$	
$E^{DF} = [] \mid (E^{DF} e^{DF}) \mid (v^{DF} e^{DF})$	(Evaluation Contexts)
$\mid (p^2 E^{DF} e^{DF}) \mid (p^2 v^{DF} E^{DF})$	
$\mid (\text{cons } E^{DF} e^{DF}) \mid (\text{cons } v^{DF} E^{DF}) \mid (p^1 E^{DF})$	
$\mid (\text{if } E^{DF} e^{DF} e^{DF})$	
$\mid (\text{force } E^{DF}) \mid (\text{force } \ell E^{DF})$	
$v^{DF} = n \mid s \mid b \mid \lambda x. e^{DF} \mid (\text{cons } v^{DF} v^{DF}) \mid \text{null} \mid \ell$	(Values)

The store in a machine configuration is represented as a list of pairs. In the above specification, ellipses means “zero or more of the preceding element”. The evaluation contexts are the standard call-by-value contexts, plus two **force** contexts. The first **force** context resembles the **force** expression in a program and indicates that some arbitrary expression is being forced. For the evaluation of a specific delayed computation, the second **force** context is used. It remembers a location so the store can be updated after the evaluation is complete. Evaluation under a $(\text{force } \ell \mid \mid)$ context corresponds to evaluation under a label in λ_{LR} .

The starting machine configuration for a Racket program e^{RKT} is $\langle e^{RKT}, () \rangle$ where the program, in an empty context, is set as the initial control string, and the store is initially empty. Evaluation stops when the control string is a value. Values are numbers, strings, booleans, abstractions, lists, or store locations. Our CS machine transitions are in figure 8. Every program e^{RKT} has a deterministic transition sequence because the left hand sides of all the transition rules are mutually exclusive and cover all possible control strings in the C register.

The β_v , PRIM, FIRST, REST, IF-TRUE, and IF-FALSE transitions are standard call-by-value transitions. The DELAY transition reduces a **delay** expression to an unused location ℓ . The delayed computation is saved at that location in the store. When the argument to a **force** expression is a location, the suspended expression at that location is retrieved from the store and plugged into a special **force** evaluation context, as specified by the FORCE-DELAY transition. The outer **force** evaluation context is retained in case there are nested delays. The special context saves the store location of the forced expression, so the store can be updated with the resulting value, as dictated by the FORCE-UPDATE transition. Finally, the FORCE-NONDELAY transition specifies that forcing a non-location value results in the removal of the outer **force** context.

4.2 Continuation Marks

A stepper for a functional language needs access to the control stack of its evaluator in order to reconstruct the evaluation steps. In a low-level stepper implementation, the stepper would be granted complete, privileged access to the control stack. As Clements et al.

$\vdash \rightarrow_{\text{CSKM}}$		
$\langle (\text{wcm } c_1^{\text{CM}} c_2^{\text{CM}}), \sigma, K^{\text{CM}}, m \rangle$	$\xrightarrow{\emptyset}_{\text{CSKM}}$	$\langle c_1^{\text{CM}}, \sigma, (\text{wcm } c_2^{\text{CM}} K^{\text{CM}} m), \emptyset \rangle$ WCM:EXP
$\langle v^{\text{CM}}, \sigma, (\text{wcm } c^{\text{CM}} K^{\text{CM}} m), m' \rangle$	$\xrightarrow{\emptyset}_{\text{CSKM}}$	$\langle c^{\text{CM}}, \sigma, K^{\text{CM}}, v^{\text{CM}} \rangle$ WCM:VAL
$\langle (\text{ccm}), \sigma, K^{\text{CM}}, m \rangle$	$\xrightarrow{\emptyset}_{\text{CSKM}}$	$\langle v^{\text{CM}}, \sigma, K^{\text{CM}}, \emptyset \rangle$ CCM where $v^{\text{CM}} = \pi(K^{\text{CM}}, m)$
$\langle (\text{output } c^{\text{CM}}), \sigma, K^{\text{CM}}, m \rangle$	$\xrightarrow{\emptyset}_{\text{CSKM}}$	$\langle c^{\text{CM}}, \sigma, (\text{output } K^{\text{CM}} m), \emptyset \rangle$ OUTPUT:EXP
$\langle v^{\text{CM}}, \sigma, (\text{output } K^{\text{CM}} m), m' \rangle$	$\xrightarrow{v^{\text{CM}}}_{\text{CSKM}}$	$\langle 42, \sigma, K^{\text{CM}}, m \rangle$ OUTPUT:VAL
$\langle (\text{loc? } c^{\text{CM}}), \sigma, K^{\text{CM}}, m \rangle$	$\xrightarrow{\emptyset}_{\text{CSKM}}$	$\langle c^{\text{CM}}, \sigma, (\text{loc? } K^{\text{CM}} m), \emptyset \rangle$ LOC?:EXP
$\langle \ell, \sigma, (\text{loc? } K^{\text{CM}} m), m' \rangle$	$\xrightarrow{\emptyset}_{\text{CSKM}}$	$\langle \text{true}, \sigma, K^{\text{CM}}, m \rangle$ LOC?-TRUE:VAL
$\langle v^{\text{CM}}, \sigma, (\text{loc? } K^{\text{CM}} m), m' \rangle$ $v^{\text{CM}} \notin \text{Locations}$	$\xrightarrow{\emptyset}_{\text{CSKM}}$	$\langle \text{false}, \sigma, K^{\text{CM}}, m \rangle$ LOC?-TRUE:VAL

Figure 9. CSKM Machine Transitions

(2001) argued, however, such privileged access is unnecessary and often undesirable.

Continuation marks are a lightweight stack-access mechanism. The stepper for Lazy Racket reuses Clements’s stepper for Racket, which utilizes continuation marks to reconstruct a program’s control stack, i.e., the evaluation context. There are two available operations for these novel values:

1. *store* a value in the current frame of the control stack,
2. *retrieve* all stored continuation marks.

Using these two operations it is possible to implement a stepper without coupling it directly to the evaluator.

Clements et al. (2001) present such a stepper for the eager Racket evaluator. The stepper first annotates a source program with continuation mark *store* and *retrieve* operations at appropriate points. Then, at each *retrieve* point, the stepper reconstructs and outputs a reduction step from information stored in the continuation marks. Our stepper extends Clements’s model with *delay* and *force* constructs. The annotation and reconstruction functions are formally defined in section 5.

4.3 Racket + delay/force + Continuation Marks

After a Lazy Racket program is expanded to a plain Racket program, the lazy stepper annotates the plain Racket program with continuation mark operations. The language λ_{CM} extends λ_{DF} in a stratified manner and models the language for annotated programs.

$$e^{\text{CM}} = e^{\text{RKT}} \mid (\text{ccm}) \mid (\text{wcm } e^{\text{CM}} e^{\text{CM}}) \mid (\text{output } e^{\text{CM}}) \mid (\text{loc? } e^{\text{CM}})$$

λ_{CM} adds four additional kinds of expressions to λ_{DF} : *wcm*, *ccm*, *output*, and a *loc?* predicate. When a *wcm*, or “with continuation mark”, expression is evaluated, its first argument is evaluated and stored in the current stack frame before its second argument is evaluated. A *ccm*, or “current continuation marks”, expression evaluates to a list of all continuation marks currently stored in the stack. When reducing an *output* expression, its argument is evaluated and sent to an output channel. An *output* expression is evaluated only for this side effect, so the result of its evaluation is thus inconsequential. The *loc?* predicate identifies locations and is needed by annotated programs.

4.4 CSKM Machine

To model continuation marks, having an explicit control stack is helpful, so we convert our CS machine to a CSK machine, where the evaluation context is separated and removed from the control string in the C register and relocated to a new K register. The conversion to a CSK machine is straightforward and is accomplished

using known techniques (Felleisen et al. 2009). In addition, we pair each context with a continuation mark m , which is stored in a fourth “M” register, giving us a CSKM machine.

For the control stack in the K register, we use an “inverted” evaluation context structure, meaning that the innermost context is now most easily accessible, giving us a more realistic stack structure. This new representation is called a continuation and there is a one-to-one correspondence between evaluation contexts and continuations. For example, an evaluation context $E^{\text{CM}}[[[\] c^{\text{CM}}]]$ becomes $(\text{app1 } c^{\text{CM}} K^{\text{CM}} m)$, where the E^{CM} context corresponds to the K^{CM} continuation. The other evaluation contexts are similarly converted. Note the extra continuation mark m associated with the continuation. Here are all the continuations K^{CM} :

$K^{\text{CM}} =$

$$\begin{aligned} & \text{mt} \mid (\text{app1 } c^{\text{CM}} K^{\text{CM}} m) \mid (\text{app2 } v^{\text{CM}} K^{\text{CM}} m) \\ & \mid (\text{prim2-1 } p2 c^{\text{CM}} K^{\text{CM}} m) \mid (\text{prim2-2 } p2 v^{\text{CM}} K^{\text{CM}} m) \\ & \mid (\text{cons1 } c^{\text{CM}} K^{\text{CM}} m) \mid (\text{cons2 } v^{\text{CM}} K^{\text{CM}} m) \\ & \mid (\text{prim1 } p1 K^{\text{CM}} m) \mid (\text{if } c_1^{\text{CM}} c_2^{\text{CM}} K^{\text{CM}} m) \\ & \mid (\text{force } K^{\text{CM}} m) \mid (\text{force } \ell K^{\text{CM}} m) \\ & \mid (\text{wcm } c^{\text{CM}} K^{\text{CM}} m) \mid (\text{output } K^{\text{CM}} m) \mid (\text{loc? } K^{\text{CM}} m) \end{aligned}$$

The configurations of the CSKM machine are:

$$\mathcal{P}^{\text{CM}} = \langle c^{\text{CM}}, \sigma, K^{\text{CM}}, m \rangle \quad (\text{Machine States})$$

$$\mathcal{S}^{\text{CM}} = \mathcal{P}_1^{\text{CM}}, \dots, \mathcal{P}_n^{\text{CM}} \quad (\text{Transition Sequences})$$

$$c^{\text{CM}} = e^{\text{CM}} \mid \ell \quad (\text{Control Strings})$$

$$v^{\text{CM}} = v^{\text{DF}} \quad (\text{Values})$$

$$\sigma = ((\ell, c^{\text{CM}}), \dots) \quad (\text{Stores})$$

$$m = \emptyset \mid v^{\text{CM}} \quad (\text{Mark Register})$$

Control strings are again extended to include location expressions, values are the same as CS machine values, and stores map locations to control string expressions. The mark register m is either empty or contains a value. For simplicity, we assume that only one mark can be associated with a continuation frame.

The transitions for our CSKM machine are in figure 9. In order to formally model output, we add an extra tag to each transition, so our machine operates as a labeled transition system (Keller 1976). When the machine emits output, the transition is tagged with the outputted value; otherwise, the transition tag is \emptyset . If a transition has no output tag, it means the output is inconsequential in the current context. In our machine, only an *output* expression emits output. For space reasons, we only include the transitions for the new

constructs: `wcm`, `ccm`, `output`, and `loc?`. The other transitions are easily derived from the transitions for the CS machine (Felleisen et al. 2009).

The starting configuration for a program e^{CM} is $\langle e^{\text{CM}}, (), \text{mt}, \emptyset \rangle$ and evaluation halts when the control string is a value and the control stack is `mt`. Every program e^{CM} has a deterministic transition sequence because the left hand sides of all the transition rules are mutually exclusive and cover all possible C and K register combinations.

The `WCM:EXP` transition sets the first argument as the control string and saves the second argument in a `wcm` continuation. In the resulting machine configuration, the m register is reinitialized to \emptyset because a new frame is pushed onto the stack. When evaluation of the first `wcm` argument is complete, the resulting value is set as the new continuation mark, as dictated by the `WCM:VAL` transition. This new continuation mark overwrites any previous mark.

The `CCM` transition uses the π function to retrieve all continuation marks in the stack. The π function is defined as follows:

$$\begin{aligned} (\pi \text{ mt } m) &= (\text{cons } m \text{ null}) \\ (\pi (\text{app1 } c^{\text{CM}} K^{\text{CM}} m) m') &= (\text{cons } m' (\pi K^{\text{CM}} m)) \\ (\pi (\text{app2 } v^{\text{CM}} K^{\text{CM}} m) m') &= (\text{cons } m' (\pi K^{\text{CM}} m)) \\ &\dots \\ (\pi (\text{loc? } K^{\text{CM}} m) m') &= (\text{cons } m' (\pi K^{\text{CM}} m)) \end{aligned}$$

Only the first few cases are shown. The rest of the definition, for other continuations, is similarly defined. The `OUTPUT:EXP` transition sets the argument in an `output` expression as the control string and pushes a new `output` continuation frame onto the control stack. Again, the continuation mark register is initialized to \emptyset due to the new stack frame. When the argument is evaluated, the resulting value is emitted as output, as modeled by the label on the `OUTPUT:VAL` transition. Finally, the `LOC?:EXP`, `LOC?:TRUE:VAL`, and `LOC?:FALSE:VAL` transitions are defined in the expected manner, producing `true` if the `loc?` predicate is applied to a location, and `false` otherwise.

5. Correctness

Unlike most IDE tools, an algebraic stepper comes with a concise specification: the λ_{LR} rewriting system. Specifically, the stepper displays λ_{LR} rewriting sequences after removing all labels from the terms. It is therefore relatively straightforward to state a correctness theorem for the stepper, assuming a function ζ that strips an λ_{LR} program of its labels.

Theorem 1 (Stepper Correctness). *If the stepper displays the sequence $e, \zeta[e_1^{\text{LR}}], \dots, \zeta[e_n^{\text{LR}}]$ for some Lazy Racket program e , then $e \mapsto_{\text{LR}} e_1^{\text{LR}} \mapsto_{\text{LR}} \dots \mapsto_{\text{LR}} e_n^{\text{LR}}$.*

The statement of the theorem's conclusion involves multistep rewriting because some rewriting steps, such as `CONS`, merely add labels and change nothing else about the term.

The proof of the theorem consists of two distinct steps. First, we show that the output of a macro-expanded, annotated Lazy Racket program uniquely describes the execution of a macro-expanded Lazy Racket program without annotations. That is, we retrieve a machine reduction sequence at the level of Racket with `delay` and `force`. Second, we prove that this reduction sequence is equivalent to the rewriting sequence of the original Lazy Racket program, modulo label assignment. The following two subsections spell out the two lemmas and present proof sketches. The proof of theorem 1 combines the two main lemmas from these subsections in a straightforward fashion.

$$\begin{aligned} \mathcal{A} : e^{\text{RKT}} &\rightarrow e^{\text{CM}} \\ \mathcal{A}[(\text{force } e^{\text{RKT}})] &= \\ &(\text{let*} \\ &([v0 \text{ (wcm (list "force") } \mathcal{A}[e^{\text{RKT}}])]) \\ &[v1 \\ &(\text{if (not (loc? v0))} \\ &v0 \\ &(\text{wcm} \\ &(\boxed{\text{(list "force")}}_1) \\ &(\text{wcm} \\ &(\boxed{\text{(list "force" v0)}}_2) \\ &(\text{let*} \\ &([v2 \text{ (force v0)}] ; v0 \text{ is location} \\ &[t0 \\ &(\boxed{\text{output}}_3 \\ &(\text{cons (list "val" } \boxed{v0}_4 \text{ } \mathcal{Q}[v2]) \\ &(\boxed{\text{(rest (ccm))}}_5))]) \\ &v2)))] \\ &[t1 \text{ (output (cons } \mathcal{Q}[v1] \text{ (ccm))})] \\ &v1) \\ \mathcal{A}[(\text{delay } e^{\text{RKT}})] &= \\ &(\text{let*} \\ &([t0 \text{ (output (cons } \mathcal{Q}[(\text{delay } e^{\text{RKT}}]) \text{ (ccm))})] \\ &[\ell \text{ (alloc)}] \\ &[t1 \text{ (output} \\ &(\text{cons (list "loc" } \ell \text{ } \boxed{\mathcal{Q}[e^{\text{RKT}}]}_6) \text{ (ccm))})] \\ &(\text{delay } \mathcal{A}[e^{\text{RKT}}]) \end{aligned}$$

Figure 10. Annotation function for `delay` and `force`.

5.1 Annotation and Reconstruction Correctness

To state the correctness lemma for the CSKM machine, we need two functions. First $\mathcal{T} : \mathcal{S}^{\text{CM}} \rightarrow v_1^{\text{CM}}, \dots, v_n^{\text{CM}}$ consumes a series of CSKM steps and produces the trace of output values:

$$\begin{aligned} \mathcal{T}[\dots \mapsto_{\text{CSKM}} \mathcal{P}_i^{\text{CM}} \xrightarrow{v^{\text{CM}}} \mathcal{P}_{i+1}^{\text{DF}} \mapsto_{\text{CSKM}} \dots] \\ = \dots, v^{\text{CM}}, \dots \end{aligned}$$

Second, $\mathcal{A} : e^{\text{RKT}} \rightarrow e^{\text{CM}}$ annotates a plain Racket program and $\mathcal{R} : v_1^{\text{CM}}, \dots, v_n^{\text{CM}} \rightarrow \mathcal{S}^{\text{DF}}$ reconstructs a CS machine transition sequence for a Racket program.

Lemma 2 (Annotation/Reconstruction Correctness). *For any Racket program e^{RKT} , if $\langle e^{\text{RKT}}, () \rangle \mapsto_{\text{CS}} \dots \mapsto_{\text{CS}} \mathcal{P}^{\text{DF}}$, then*

$$\begin{aligned} \mathcal{R}[\mathcal{T}[\langle \mathcal{A}[e^{\text{RKT}}], () \rangle \mapsto_{\text{CSKM}} \dots \mapsto_{\text{CSKM}} \mathcal{P}^{\text{CM}}]] \\ = \langle e^{\text{RKT}}, () \rangle \mapsto_{\text{CS}} \dots \mapsto_{\text{CS}} \mathcal{P}^{\text{DF}} \end{aligned}$$

Our annotation and reconstruction functions extend the functions of Clements et al. (2001). Figures 10 and 11 summarize these additions. We omit the parts defined by Clements et al. and instead review the functions with some illustrative examples.

Annotation adds `output` expressions and continuation mark `wcm` and `ccm` operations to a program. For example, annotating the program `(+ 1 2)` results in the following annotated program:⁵

⁵For clarity, the syntactic sugar `let*` and `list` forms are used. They are defined in the standard way. Other minor code-readability improvements have also been made.

$$\begin{aligned}
& \mathcal{R} : v_1^{\text{cm}}, \dots, v_n^{\text{cm}} \rightarrow \mathcal{S}^{\text{DF}} \\
& \mathcal{R}[\dots, v_i^{\text{cm}}, \dots] = \dots, \langle \mathcal{R}_E[(\text{rest } v_i^{\text{cm}})] [\mathcal{R}_C[(\text{first } v_i^{\text{cm}})]], \\
& \quad \mathcal{R}_S[v_1^{\text{cm}}, \dots, v_i^{\text{cm}}], \\
& \quad \dots \\
& \quad \mathcal{R}_E : v^{\text{cm}} \rightarrow E^{\text{DF}} \\
& \quad \mathcal{R}_E[(\text{cons } (\text{list "force"}) v^{\text{cm}})] = (\text{force } \mathcal{R}_E[v^{\text{cm}}]) \\
& \quad \mathcal{R}_E[(\text{cons } (\text{list "force"} \ell) v^{\text{cm}})] = (\text{force } \ell \mathcal{R}_E[v^{\text{cm}}]) \\
& \quad \mathcal{R}_C : v^{\text{cm}} \rightarrow e^{\text{DF}} \\
& \quad \mathcal{R}_C[(\text{list "val"} \ell v^{\text{DF}})] = \mathcal{Q}^{-1}[v^{\text{DF}}] \\
& \quad \mathcal{R}_C[(\text{list "loc"} \ell v^{\text{DF}})] = \ell \\
& \quad \text{otherwise, } \mathcal{R}_C[v^{\text{DF}}] = \mathcal{Q}^{-1}[v^{\text{DF}}] \\
& \quad \mathcal{R}_S : v_1^{\text{cm}}, \dots, v_n^{\text{cm}} \rightarrow \sigma \\
& \quad \mathcal{R}_S[(\text{cons } (\text{list } s \ell v^{\text{DF}'}) v^{\text{DF}'}, v_{\text{rest}}^{\text{DF}}, \dots)] = \\
& \quad \quad \begin{cases} (\text{cons } (\ell, \mathcal{Q}^{-1}[v^{\text{DF}'}]) \mathcal{R}_S[v_{\text{rest}}^{\text{DF}}, \dots]) \\ \quad \text{if } \ell \notin \text{dom}(\mathcal{R}_S[v_{\text{rest}}^{\text{DF}}, \dots]) \\ \mathcal{R}_S[v_{\text{rest}}^{\text{DF}}, \dots] \\ \quad \text{if } \ell \in \text{dom}(\mathcal{R}_S[v_{\text{rest}}^{\text{DF}}, \dots]) \end{cases} \\
& \quad \mathcal{R}_S[v^{\text{DF}}, v_{\text{rest}}^{\text{DF}}, \dots] = \mathcal{R}_S[v_{\text{rest}}^{\text{DF}}, \dots] \\
& \quad (\text{first } v^{\text{DF}}) \neq (\text{list } s \ell v^{\text{DF}'})
\end{aligned}$$

Figure 11. Reconstruction function for delay and force.

```

(let*
  ([t0 (output (cons Q[(+ 1 2)] (ccm)))]
   [v1 (+ 1 2)]
   [t1 (output (cons Q[v1] (ccm)))]
  v1)

```

Annotated programs utilize the quoting function \mathcal{Q} , which converts an expression into a value representation. For example $\mathcal{Q}[(+ 1 2)] = (\text{list "+" 1 2})$. There is also an inverse function, \mathcal{Q}^{-1} , for reconstruction. The above annotated program evaluates to 3, outputting the values $\mathcal{Q}[(+ 1 2)]$ and $\mathcal{Q}[3]$ in the process, from which the reduction sequence $(+ 1 2) \rightarrow 3$ can be recovered. The (ccm) calls in the example return the empty list since no continuation marks were previously stored, i.e., there were no calls to wcm. There is no need for wcm annotations because the entire program is a redex, i.e., the context is empty.

Extending the example to $(+ (+ 1 2) 5)$ yields

```

(let*
  ([v0 (wcm (list "prim2-1" "+" 5))
   (let*
     ([t0 (output (cons Q[(+ 1 2)] (ccm)))]
      [v1 (+ 1 2)]
      [t1 (output (cons Q[v1] (ccm)))]
     v1))]
   [v2 (+ v0 5)]
   [t2 (output (cons Q[v2] (ccm)))]
  v2)

```

This extended example contains the first program as a subexpression; and therefore the annotated version of the program contains the annotated version of the first example. The $(+ 1 2)$ expression now occurs in the context $(+ [] 5)$ and the wcm expression

stores an appropriate continuation mark so this context can be reconstructed. The "prim2-1" label indicates that the hole is in the left argument position. The first output expression now produces the output value $(\text{list } \mathcal{Q}[(+ 1 2)] (\text{list "prim2-1" "+" 5}))$, which can be reconstructed to the expression $(+ (+ 1 2) 5)$. Reconstructing all outputs produces $(+ (+ 1 2) 5) \rightarrow (+ 3 5) \rightarrow 8$.

Storing context information in continuation marks also enables the reconstruction of a machine state, which is what the stepper actually does, and the reconstruction and annotation functions defined in figures 10 and 11 demonstrate how this works. Figure 10 shows that if the subexpression e^{RKT} of a force expression does not evaluate to location, the annotations are like those for the above examples. If e^{RKT} produces a location, additional continuation marks (figure 10, boxes 1 and 2) are needed to indicate the presence of force contexts during evaluation of a delayed computation. An additional output expression (box 3) is also needed so that the steps showing the removal of both the (force []) and (force ℓ []) contexts can be reconstructed. Note the (rest (ccm)) (box 5) in the first output; this ensures the (force ℓ []) context is not part of the reconstructed control string. The location $v0$ (box 4) is included in the output so the store can be properly reconstructed. The "val" tag directs the reconstruction function to use the value $\mathcal{Q}[v2]$ from the emitted location-value pair for reconstructing the control string.

The annotation of a delay expression requires predicting the location of the delayed computation in the store. We therefore assume we have access to an alloc function that uses the same location-allocating algorithm as the memory management system of the machine.⁶ In addition to the location, the delayed expression itself (box 6) is also included in the output, to enable reconstruction of the store. The "loc" tag directs the reconstruction function to use the location from the emitted location-value pair for reconstructing the control string.

The reconstruction function in figure 11 consumes a list of values, where each value is a sublist, and reconstructs a CS machine state from each sublist. Again, only the cases involving delay and force are defined. The rest of the function is borrowed from Clements (2006). The first element of every v_i^{cm} sublist represents a (quoted) expression that is plugged into the context represented by the rest of the sublist. The store is reconstructed by retrieving all the location-value pairs in all the sublists up to the current one. The arguments to the store-reconstruction function \mathcal{R}_S may contain duplicate entries for a location, so a location-value pair is only included in the resulting store if it does not occur in any subsequent arguments.

Lemma 2 Proof Sketch. The proof of lemma 2 extends Clements's proof (Clements 2006, Section 3.4) with cases for delay and force. Also we modify the argumentation for the original cases to cope with the additional store, but doing so is straightforward because the store is simply threaded through. \square

5.2 Lazy Racket Correctness

The function $\varphi : e \rightarrow e^{\text{RKT}}$ macro-expands a Lazy Racket program. Because source terms don't include labels, φ is undefined for labeled terms. Its partial inverse $\varphi^{-1} : e^{\text{DF}} \times \sigma \rightarrow e$ synthesizes an unlabeled Lazy Racket program from a (CS machine representation of a) plain Racket program. The expansion and synthesis functions are defined in figures 12 and 13, respectively.

Lemma 3 states the correctness of Lazy Racket in terms of the functions φ , φ^{-1} , ζ , and the λ_{LR} rewriting system from section 3. That is, every CS machine transition sequence has an equivalent λ_{LR} rewriting sequence, modulo φ^{-1} and ζ .

⁶Since labels are not displayed as numbers but as sharing among expressions, this unrealistic mathematical assumption is acceptable.

$$\begin{aligned}
\varphi : e &\rightarrow e^{\text{RKT}} \\
\varphi[\lambda x. e] &= \lambda x. \varphi[e] \\
\varphi[(e_1 \ e_2)] &= ((\text{force } \varphi[e_1]) (\text{delay } \varphi[e_2])) \\
\varphi[(p^2 \ e_1 \ e_2)] &= (p^2 (\text{force } \varphi[e_1]) (\text{force } \varphi[e_2])) \\
\varphi[(\text{cons } e_1 \ e_2)] &= (\text{cons } (\text{delay } \varphi[e_1]) (\text{delay } \varphi[e_2])) \\
\varphi[(p^1 \ e)] &= (p^1 (\text{force } \varphi[e])) \\
\varphi[(\text{if } e_1 \ e_2 \ e_3)] &= (\text{if } (\text{force } \varphi[e_1]) \varphi[e_2] \varphi[e_3]) \\
\text{otherwise, } \varphi[e] &= e
\end{aligned}$$

Figure 12. Macro-expanding Lazy Racket.

Lemma 3 (LR Correctness). *For all Lazy Racket programs e and Racket programs c^{DF} such that $\langle \varphi[e], () \rangle \mapsto_{\text{CS}} \langle c^{\text{DF}}, \sigma \rangle$, there exists a Lazy Racket program e^{LR} such that $e \mapsto_{\text{LR}} e^{\text{LR}}$ and $\varphi^{-1}[[c^{\text{DF}}]]\sigma = \zeta[[e^{\text{LR}}]]$.*

Proof Sketch. We prove the lemma by induction on the number of CS machine steps. For the base case, the lemma holds because $\varphi^{-1}[[\varphi[e]]](\) = e$. Otherwise, we proceed by case analysis on the last transition step.

For each case, we prove correct synthesis of evaluation contexts and redexes separately. When the rewriting of a redex affects the context, too—see figure 5, third column—we need information about the redex for the synthesis of a context. This parallel rewriting of λ_{LR} programs is equivalent to the reduction of stored expressions in the CS machine if there are multiple references to that expression. If the stored expression is a value, the reconstruction naturally reifies the value throughout because φ^{-1} translates locations by retrieving and unexpanding the expression at that store location.

If a stored expression is in the process of being reduced to a value, however, the store has not been updated but all references to this location must reflect the partial reductions. Such intermediate steps manifest themselves as reductions under $(\text{force } \ell \ [\])$ contexts in plain Racket. To synthesize these intermediate steps properly, φ^{-1} first updates the store with the partially reduced expression and then synthesizes the rest of the state into a Lazy Racket program. As a result, the translation of all occurrences of a location yield the desired intermediate expression. Only synthesis of $(\text{force } \ell \ c^{\text{DF}})$ control strings yields a new store; thus only subexpressions that possibly contain $(\text{force } \ell \ c^{\text{DF}})$, i.e., subexpressions that can contain the redex, yield a new store.

For synthesis of a β_v redex, the substituted argument is a location because φ wraps all application arguments in a delay . The φ^{-1} synthesis function translates locations to the delayed expression at that location, so the corresponding λ_{LR} step is β_{LR} . The β_{LR} step adds a label, which is removed with ζ .

For synthesis of PRIM, FIRST, REST, IF-TRUE, or IF-FALSE redexes, φ^{-1} yields a PRIM, FIRST, REST, IF-TRUE, or IF-FALSE λ_{LR} redex, respectively, and these rewriting steps insert no labels.

If the last CS step, $\langle c^{\text{DF}}, \sigma \rangle \mapsto_{\text{CS}} \langle c^{\text{DF}'}, \sigma' \rangle$, follows DELAY, FORCE-DELAY, FORCE-UPDATE, or FORCE-NONDELAY, there is no corresponding λ_{LR} rewriting rule and $\varphi^{-1}[[c^{\text{DF}}]]\sigma = \varphi^{-1}[[c^{\text{DF}'}]]\sigma'$. The lemma holds since this is equivalent to taking zero steps. \square

6. Performance

The performance of a stepper tool must be measured against the programmer's ability to use the tool. In particular, raw performance numbers are inconsequential because the usability of the tool primarily depends on its responsiveness to I/O.

$$\begin{aligned}
\varphi^{-1} : c^{\text{DF}} \times \sigma &\rightarrow e \\
\varphi^{-1}[[c^{\text{DF}}]]\sigma &= e \\
&\text{where } \langle e, \sigma' \rangle = \bar{\varphi}[[c^{\text{DF}}]]\sigma \\
\bar{\varphi} : c^{\text{DF}} \times \sigma &\rightarrow \langle e, \sigma \rangle \\
\bar{\varphi}[\lambda x. e^{\text{DF}}]\sigma &= \langle \lambda x. e, \sigma \rangle \\
&\text{where } \langle e, \sigma \rangle = \bar{\varphi}[[e^{\text{DF}}]]\sigma \\
\bar{\varphi}[(c_1^{\text{DF}} \ c_2^{\text{DF}})]\sigma &= \langle (e_1 \ e_2), \sigma'' \rangle \\
&\text{where } \langle e_1, \sigma' \rangle = \bar{\varphi}[[c_1^{\text{DF}}]]\sigma \\
&\quad \langle e_2, \sigma'' \rangle = \bar{\varphi}[[c_2^{\text{DF}}]]\sigma' \\
\bar{\varphi}[(p^2 \ c_1^{\text{DF}} \ c_2^{\text{DF}})]\sigma &= \langle (p^2 \ e_1 \ e_2), \sigma'' \rangle \\
&\text{where } \langle e_1, \sigma' \rangle = \bar{\varphi}[[c_1^{\text{DF}}]]\sigma \\
&\quad \langle e_2, \sigma'' \rangle = \bar{\varphi}[[c_2^{\text{DF}}]]\sigma' \\
\bar{\varphi}[(\text{cons } c_1^{\text{DF}} \ c_2^{\text{DF}})]\sigma &= \langle (\text{cons } e_1 \ e_2), \sigma'' \rangle \\
&\text{where } \langle e_1, \sigma' \rangle = \bar{\varphi}[[c_1^{\text{DF}}]]\sigma \\
&\quad \langle e_2, \sigma'' \rangle = \bar{\varphi}[[c_2^{\text{DF}}]]\sigma' \\
\bar{\varphi}[(p^1 \ c^{\text{DF}})]\sigma &= \langle (p^1 \ e), \sigma' \rangle \\
&\text{where } \langle e, \sigma' \rangle = \bar{\varphi}[[c^{\text{DF}}]]\sigma \\
\bar{\varphi}[(\text{if } c_1^{\text{DF}} \ c_2^{\text{DF}} \ c_3^{\text{DF}})]\sigma &= \langle (\text{if } e_1 \ e_2 \ e_3), \sigma' \rangle \\
&\text{where } \langle e_1, \sigma' \rangle = \bar{\varphi}[[c_1^{\text{DF}}]]\sigma \\
&\quad \langle e_2, \sigma' \rangle = \bar{\varphi}[[c_2^{\text{DF}}]]\sigma' \\
&\quad \langle e_3, \sigma' \rangle = \bar{\varphi}[[c_3^{\text{DF}}]]\sigma' \\
\bar{\varphi}[(\text{delay } e^{\text{DF}})]\sigma &= \bar{\varphi}[[e^{\text{DF}}]]\sigma \\
\bar{\varphi}[\ell]\sigma &= \bar{\varphi}[[\sigma[\ell]]]\sigma \\
\bar{\varphi}[(\text{force } c^{\text{DF}})]\sigma &= \bar{\varphi}[[c^{\text{DF}}]]\sigma \\
\bar{\varphi}[(\text{force } \ell \ c^{\text{DF}})]\sigma &= \langle e, \sigma'[\ell \leftarrow e] \rangle \\
&\text{where } \langle e, \sigma' \rangle = \bar{\varphi}[[c^{\text{DF}}]]\sigma \\
\text{otherwise, } \bar{\varphi}[[e^{\text{DF}}]]\sigma &= \langle e^{\text{DF}}, \sigma \rangle
\end{aligned}$$

Figure 13. Synthesizing Lazy Racket from plain Racket.

Nevertheless, it is important to quantify the basic performance penalty of a stepper. In our case, the stepper slows down the execution of programs by a factor of 20 (up to 60), as the following table for a small number of micro-benchmarks shows:

Test Name	Slowdown Factor
fibo	21.4
ack	32.7
partial	27.3
tak	23.0
takl	34.9
takr	55.5

The numbers describe the performance ratio of annotated versus unannotated programs. While they include both annotation and evaluation time, the annotation time is negligible when compared to the time it takes to run the program.

We achieve adequate responsiveness of the stepper with a straightforward arrangement. As soon as the stepper backend produces output, the front-end asynchronously displays reduction

steps. By the time the programmer has read and understood the first step, the stepper can quickly respond to additional requests.

7. Related Work

Debugging programs in lazy languages poses serious challenges, as numerous attempts have shown over the past three decades. Most recent debuggers are for Haskell. Allwood et al. (2009) developed a stack-trace-generating tool for GHC that aids in identifying the context of an error. Their tool transforms programs to carry around an extra stack parameter for accessing the program context, similar to continuation marks. Their tool approximates a call-by-value stack, however, instead of displaying the lazy evaluation order.

Ennals and Peyton Jones (2003a) built HsDebug for GHC, which employs the “stop, examine, continue” style found in imperative languages, e.g., gdb. With HsDebug, programmers can set breakpoints and can examine the program state at breakpoints. Like Allwood’s debugger, HsDebug does not preserve laziness in programs; the debugger must make certain tweaks, resulting in an “optimistic” evaluation model (Ennals and Peyton Jones 2003b).

The most recent version of the GHC system has debugging features built into its interpreter (Marlow et al. 2007). In principle, the GHCi debugger can most closely mimic the operation of our stepper because: (1) it allows a programmer to single-step through the evaluation of a program, and (2) it preserves laziness for the user to observe. In contrast to our stepper, the GHCi debugger does not use a substitution semantics, and only displays a few lines of the program at a time, resulting in frequent jumps from the body of a function to different call sites, when arguments need to be evaluated. This jumping can be confusing for a programmer to follow, especially a novice to lazy evaluation. Also, sharing is not easily observed in the GHCi debugger. All thunks are rendered the same way (as a “.” character), and when stepping, the debugger skips over the reduction of a variable, so it is occasionally unclear which argument is being evaluated.

Hat 2.0 (Chitil et al. 2003; Wallace et al. 2001) is a suite of debugging tools for GHC that aggregates and improves on several previous tools: Hat 1.0 (Sparud and Runciman 1997), Hood (Gill 2000), and Freja (Nilsson 1998). Hat’s implementation resembles the implementation of our stepper in that it works by transforming a source program into an annotated one, which, when run, generates trace information. The generated trace is then interpreted by the tools. In Hat, the generated trace can be viewed in several different styles: as a directed graph of expressions connected according to redex-contractum relations (like the original Hat), in the question-answer style of an algorithmic debugger (like Freja), or by tracking specific values in a computation (like Hood). The graph viewer is somewhat similar to our tool; however, program evaluation is portrayed using graph semantics, and when viewing the graph, it can be difficult to deduce the order of reductions for graphs larger than a few nodes and links. Hat also includes a viewer, hat-stack, that shows a simulated eager evaluation stack, similar to Allwood’s tool. Hat is no longer maintained and it is not clear if it still works with the more recent version of GHC.

Finally, only one of these tools come with formal models and correctness proofs for their architecture. Chitil and Luo (Chitil and Luo 2006) developed a model for Hat’s trace generator and show that the evaluation steps can be reconstructed from the information in the traces.

8. Conclusion

We have presented a lazy stepper as an additional tool for the lazy language debugging arsenal. The stepper presents computation as the standard rewriting sequence of a novel lazy semantics. While the stepper is implemented in Lazy Racket via a “macro” over the

existing stepper for strict Racket, our paper explains the general software architecture via a generic theoretical model. We conjecture that it is straightforward to construct a stepper on top of other architectures, like Allwood’s StackTrace.

References

- T. O. Allwood, S. Peyton Jones, and S. Eisenbach. Finding the needle: stack traces for GHC. In *Proc. 2nd Symp. on Haskell*, pages 129–140, 2009.
- Z. M. Ariola and M. Felleisen. The call-by-need lambda calculus. *J. of Functional Programming*, 7(3):265–301, 1997.
- H. P. Barendregt. *The Lambda Calculus: Its Syntax and Semantics*. North Holland, 1984.
- E. Barzilay and J. Clements. Laziness without all the hard work. In *Proc. Works. on Functional and Declarative Programming in Education*, pages 9–13, 2005.
- O. Chitil and Y. Luo. Structure and properties of traces for functional programs. In *Proc. 3rd International Works. on Term Graph Rewriting*, pages 39–63, 2006.
- O. Chitil, C. Runciman, and M. Wallace. Transforming haskell for tracing. In *Revised Selected Papers 14th International Works. on the Implementation of Functional Languages*, pages 165–181, 2003.
- J. Clements. *Portable and High-level Access to the Stack with Continuation Marks*. PhD thesis, Northeastern University, 2006.
- J. Clements, M. Flatt, and M. Felleisen. Modeling an algebraic stepper. In *Proc. 10th European Symp. on Programming*, pages 320–334, 2001.
- R. Ennals and S. Peyton Jones. Hsdebug: debugging lazy programs by not being lazy. In *Proc. Works. on Haskell*, pages 84–87, 2003a.
- R. Ennals and S. Peyton Jones. Optimistic evaluation: an adaptive evaluation strategy for non-strict programs. In *Proc. 8th International Conf. on Functional Programming*, pages 287–298, 2003b.
- M. Felleisen, R. B. Findler, and M. Flatt. *Semantics Engineering with PLT Redex*. MIT Press, 2009.
- R. B. Findler, J. Clements, C. Flanagan, M. Flatt, S. Krishnamurthi, P. Steckler, and M. Felleisen. DrScheme: A programming environment for Scheme. *J. of Functional Programming*, 12(2):159–182, 2002.
- D. P. Friedman and D. S. Wise. Cons should not evaluate its arguments. In *Proc. 3rd International Colloq. on Automata, Languages and Programming*, pages 256–284, 1976.
- A. Gill. Debugging haskell by observing intermediate data structures. In *Proc. Works. on Haskell*, pages 1–12, 2000.
- J. Hatcliff and O. Danvy. Thunks and the λ -calculus. *J. of Functional Programming*, 7(3):303–319, 1997.
- J. Hughes. Why functional programming matters. *Computer J.*, 32(2):98–107, 1989.
- R. M. Keller. Formal verification of parallel programs. *Communications of the ACM*, 19(7):371–384, 1976.
- J. Launchbury. A natural semantics for lazy evaluation. In *Proc. 20th Symp. on Principles of Programming Languages*, pages 144–154, 1993.
- J. Maraist, M. Odersky, and P. Wadler. The call-by-need lambda calculus. *J. of Functional Programming*, 8(3):275–317, 1998.
- S. Marlow, J. Iborra, B. Pope, and A. Gill. A lightweight interactive debugger for haskell. In *Proc. Works. on Haskell*, pages 13–24, 2007.
- H. Nilsson. *Declarative Debugging for Lazy Functional Languages*. PhD thesis, Linköpings Universitet, 1998.
- S. Peyton Jones, C. Hall, K. Hammond, J. Cordy, H. Kevin, W. Partain, and P. Wadler. The glasgow haskell compiler: a technical overview, 1992.
- G. D. Plotkin. Call-by-name, call-by-value, and the λ -calculus. *Theoretical Computer Science*, 1(2):125–159, 1975.
- J. Sparud and C. Runciman. Tracing lazy functional computations using redex trails. In *Proc. 9th International Symp. on Programming Languages, Implementations, Logics and Programs*, pages 291–308, 1997.
- M. Wallace, O. Chitil, T. Brehm, and C. Runciman. Multiple-view tracing for haskell: a new hat. In *Proc. Works. on Haskell*, pages 151–170, 2001.